

Remarks

Allowance of all claims is respectfully requested in view of the amendments and remarks presented. Claims 1, 2, 7, 9-30, 32-54 & 56-74 remain pending.

By this paper, Applicants have further amended independent claims 1, 28, 51 & 52 to more particularly point and distinctly claim certain features of Applicants' invention. These amendments to the claims constitute a *bona fide* attempt by Applicants to advance prosecution of the application and obtain allowance of certain claims, and are in no way meant to acquiesce to the substance of the Office Action. It is believed that the amendments to the claims place all claims in condition for allowance.

In the independent claims, Applicants further characterize the local configuration data as being located at a local data storage, and the global configuration data as being located at a global data repository, wherein the global data repository is a global data repository for the cluster. Additionally, the independent claims further characterize the recited "define cluster operation" used to initially automatically create the cluster. This define cluster operation includes initially creating a new cluster identifier for the cluster, and defining the local data storage area for a node issuing the define cluster operation, and defining the global data repository area for the cluster. Support for the claim amendments can be found throughout the application as filed, for example, reference pages 14-35 of the specification. No new matter is added to the application by any amendment presented.

Claims 1, 2, 7, 9-16, 19-23, 26-30, 32-39, 42-46, 49-54, 56-63, 66-70, 73 & 74 were rejected under 35 U.S.C. 102(e) as being anticipated by Slaughter et al. (U.S. Patent No. 6,014,669), while claims 17, 18, 40, 41, 64 & 65 were rejected under 35 U.S.C. 103(a) as being unpatentable over Slaughter et al. in view of Modiri et al. (U.S. Patent No. 6,192,401), and claims 24, 25, 47, 48, 71 & 72 were rejected under U.S.C. 103(a) as being unpatentable over Slaughter et al. in view of Zhang et al. (U.S. Patent No. 5,832,182). Each of these rejections is respectfully, but most strenuously, traversed to any extent deemed applicable to the amended claims submitted herewith.

In one aspect, applicants' invention is directed to a method of managing cluster configurations of a computing environment (e.g., independent claim 1). The method includes executing a distributed configuration management component on a plurality of nodes of a cluster of the computing environment; and providing configuration consistency for the cluster and cluster membership control using the distributed configuration management component. The providing configuration consistency includes performing a comparison operation between local configuration data at a local data storage and global configuration data at a global data repository, wherein the global data repository comprises a global data repository for the cluster, while the providing cluster membership control includes providing one or more cluster membership control operations associated with the cluster. The one or more cluster membership control operations include a define cluster operation used to initially automatically create the cluster, the define cluster operation initially creating a new cluster identifier for the cluster, and defining the local data storage area for a node issuing the define cluster operation, and defining the global data repository area or for the cluster. Thus, in applicants' claimed invention, definition of the cluster is performed via an explicit define cluster operation, making the define cluster function independent of the actual hardware components of the computing environment.

With respect to the anticipation rejection, it is well settled that there is no anticipation unless (1) all the same elements are (2) found in exactly the same situation and (3) are united in the same way to (4) perform the identical function. As submitted herewith, applicants' independent claims and the Slaughter et al. patent clearly do not teach the same elements or even functional equivalents thereof. To the contrary, there are significant patentable differences between the subject matter recited in applicants' independent claims and the highly available distributed cluster configuration database teachings of Slaughter et al.

For example, applicants' independent claims each recite providing configuration consistency of the cluster using the distributed configuration management component. This providing configuration consistency includes performing a comparison between local configuration data at a local data storage and global configuration data at a global data repository. Additionally, the independent claims recite the distributed configuration management component is used to provide cluster membership control. This providing of cluster management control includes providing one or more cluster membership control operations

associated with the cluster. The one or more cluster membership control operations includes a define cluster operation that is used initially to automatically create the cluster. Thus, the distributed configuration management component executing on the one or more nodes of the computing environment is the mechanism in applicants' invention by which the cluster itself is defined. That is, there is a software command in applicants' cluster membership control functionality of the distributed configuration management component. The independent claims further specify that this define cluster operation includes initially creating a new cluster identifier for the cluster, and defining the local data storage area for a node issuing the define cluster operation, and defining the global data repository area for the cluster. The above-noted aspects of applicants' independent claims are discussed further below.

While Slaughter et al. is concerned with providing a consistent configuration database, the manner in which Slaughter et al. achieves this goal is different from that claimed by applicants given certain architectural differences in the computing environments. For example, Slaughter et al. do not use a global configuration database or repository, but only describe data stored locally at each node. Thus, there can be no comparison between local configuration data and global configuration data, as claimed by applicants. To further describe, in Slaughter et al., each node in the cluster maintains its own copy of a configuration database, and thus, configuration database operations can be performed from any node (see, e.g., Abstract, lines 4-6). This is local data to each node. The consistency of this local data is verified from a consistency record, which is located within the local data (see, e.g., col. 2, lines 45-47):

Each local copy of the configuration database uses a self-contained consistency record to uniquely identify and stamp each copy of the configuration database. The consistency of each local copy of the configuration database can be verified from the consistency record.

Thus, in Slaughter et al., the consistency record in the local data is used to determine whether the local data is invalid. There is no comparison between local configuration data at local data storage and global configuration data at the global data repository, as claimed by applicants.

Slaughter et al. explicitly chose an implementation that avoids using any global or central repository. This is described throughout the Background of Slaughter et al., in which the problems associated with such repositories are discussed. Instead of using global repositories,

Slaughter et al. use local repositories on each node. If one of the local databases is invalid, as indicated by the local consistency record, it is replaced by a database from another node. Again, there is no teaching of a comparison between local configuration data and global configuration data.

Slaughter et al. do make mention of global consistency, or cluster wide consistency, in which the cluster configuration database uses a two-phase commit protocol to guarantee the consistency of the configuration database after a configuration database update (see, e.g., col. 5, lines 30-34). This description of global consistency, however, is different from a description of comparing local configuration data with global configuration data. The global consistency in Slaughter et al. is merely ensuring that all the local copies remain consistent. Again, there is no description, teaching or suggestion of comparing global data with local data. Instead, to provide global consistency, each node in Slaughter et al. stores a backup copy of the configuration database, and if an update fails, then the node can be restored to the backup copy. There is no description, teaching or suggestion of comparing local data with global data. Thus, applicants respectfully submit that Slaughter et al. do not teach or suggest applicants' claimed invention.

To support a rejection of comparing data in local storage with data in global storage, the Office Action refers to col. 4, lines 48-52; col. 5, lines 24-27; col. 7, lines 32-36; and col. 7, lines 58 – col. 8, line 4 of Slaughter et al. (see rejection of claims 2 and 7). However, applicants respectfully submit that a careful reading of those sections, as well as the remainder of Slaughter et al., does not uncover any description or suggestion of a comparison between local configuration data and global configuration data at structures such as defined by applicants in the independent claims. For instance, col. 4, lines 48-52 merely state that a reconfiguration algorithm is provided to update a cluster configuration database and maintain consistent data. No comparison is described. Further, col. 5, lines 24-27 merely describe a two level consistency framework in which the cluster configuration database first checks local consistency and then global consistency. Again, the local consistency is determined using a local consistency record, and thus, there is no comparison between local configuration data and global configuration data. Moreover, the global consistency in Slaughter et al. is just an ensurance that the configuration database on each node is consistent after a configuration database update. Since Slaughter et al. fail to teach or suggest applicants' claimed feature of providing configuration consistency which

includes performing a comparison between local configuration data and global configuration data, applicants respectfully submit that Slaughter et al. do not anticipate applicants' claimed invention. Thus, applicants respectfully request an indication of allowability for claim 1, as well as the other independent claims.

Still further, the independent claims presented herewith recite that the distributed configuration management component is used to provide cluster membership control. More particularly, this providing of cluster membership control includes providing one or more cluster membership control operations associated with the cluster. The one or more cluster membership control operations specifically includes a define cluster operation used to initially automatically define the cluster. Applicants respectfully submit that a careful reading of Slaughter et al. fails to uncover any teaching, suggestion or implication of a distributed configuration management component executing on a plurality of nodes which provides a define cluster operation that is used to initially create a cluster. **Essentially, applicants provide an administrative command as part of the distributed configuration management component executing on the nodes which allows for the definition of a cluster of nodes.** This is in sharp contrast to Slaughter et al. wherein the cluster is defined in hardware by the physical connection of nodes to a communications interface 102 (see FIGs. 1-3 & 5 of Slaughter et al., as well as the supporting description thereof). For this reason alone, applicants respectfully submit that there is no anticipation of the independent claims presented based upon Slaughter et al.

Yet further, Applicants characterize the define cluster operation set forth in the independent claims provided herewith. As recited, the define cluster operation initially creates a new cluster identifier for the cluster, defines the local data storage area for the node issuing the define cluster operation, and defines the global data repository area for the cluster. No similar functionality is taught or suggested by Slaughter et al. Again, defining of a cluster in Slaughter et al. comprises a physical connection of nodes to a communications interface. By comparison, the present invention allows for the automatic creation of a cluster by the issuance of the define cluster operation from one node of a plurality of nodes of the computing environment.

For all the above reasons, Applicants respectfully submit that the independent claims presented herewith patentably distinguish over the teachings of Slaughter et al.

The remaining applied art, Modiri et al. and Zhang et al. fails alone or in combination to teach, suggest or imply the above-noted deficiencies of Slaughter et al. when applied against the independent claims presented. Both Modiri et al. and Zhang et al. are cited in the Office Action for various aspects of applicants' invention recited in certain dependent claims at issue. For the reasons noted above, applicants respectfully submit that the independent claims presented patentably distinguish over the applied art. The dependent claims are believed allowable for the same reasons as the independent claims from which they directly or ultimately depend, as well as for their own additional characterizations.

All pending claims are believed to be in condition for allowance and such action is respectfully requested.

Should the Examiner wish to discuss this case with applicants' attorney, the Examiner is invited to contact applicants' representative at the below-listed number.

Respectfully submitted,



Kevin P. Radigan
Attorney for Applicants
Registration No.: 31,789

Dated: June 15, 2004.

HESLIN ROTHENBERG FARLEY & MESITI P.C.
5 Columbia Circle
Albany, New York 12203-5160
Telephone: (518) 452-5600
Facsimile: (518) 452-5579